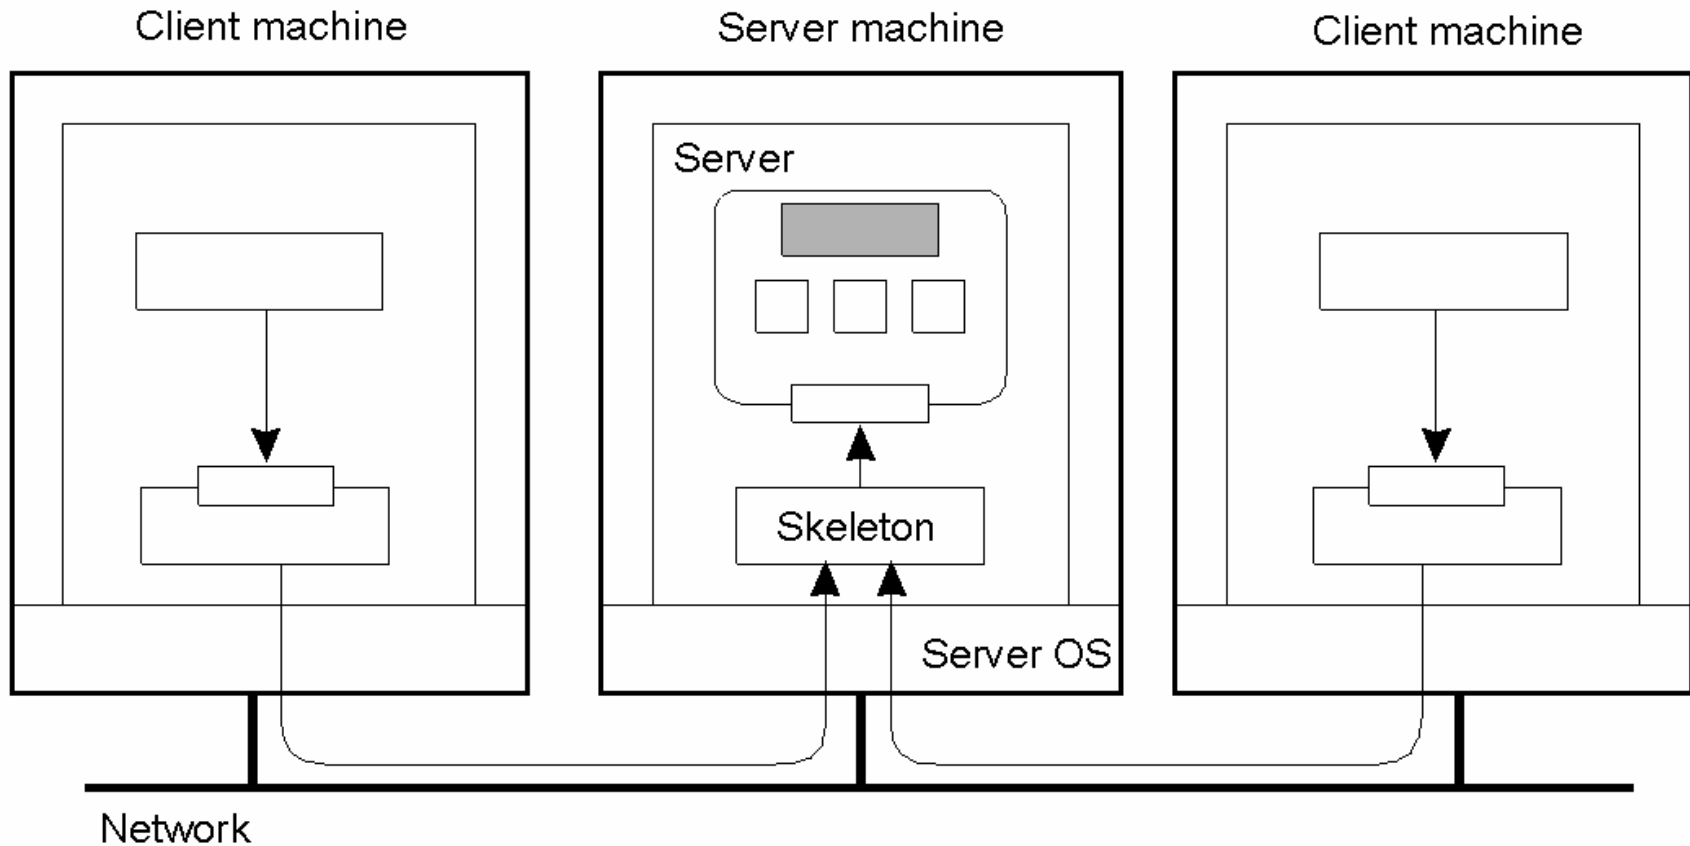


# Consistency and Replication

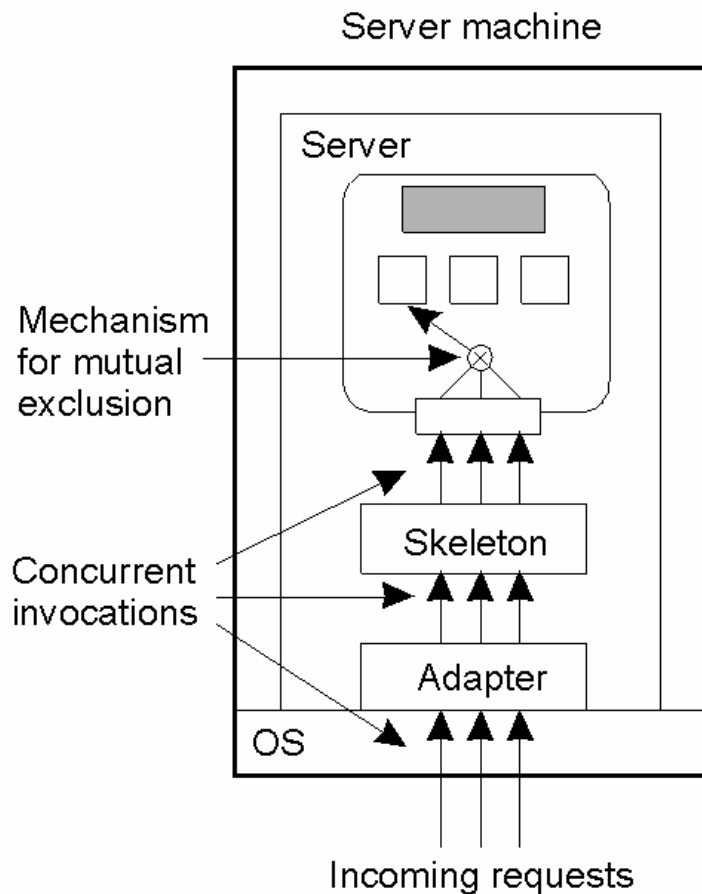
Raul Ceretta Nunes

# Object Replication (1)

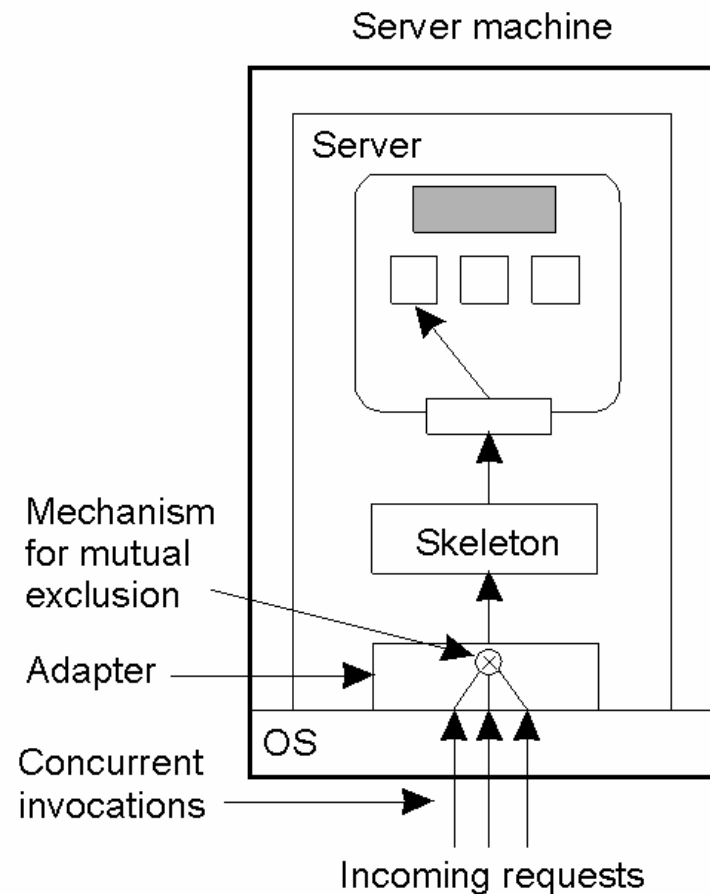


Organization of a distributed remote object shared by two different clients.

# Object Replication (2)



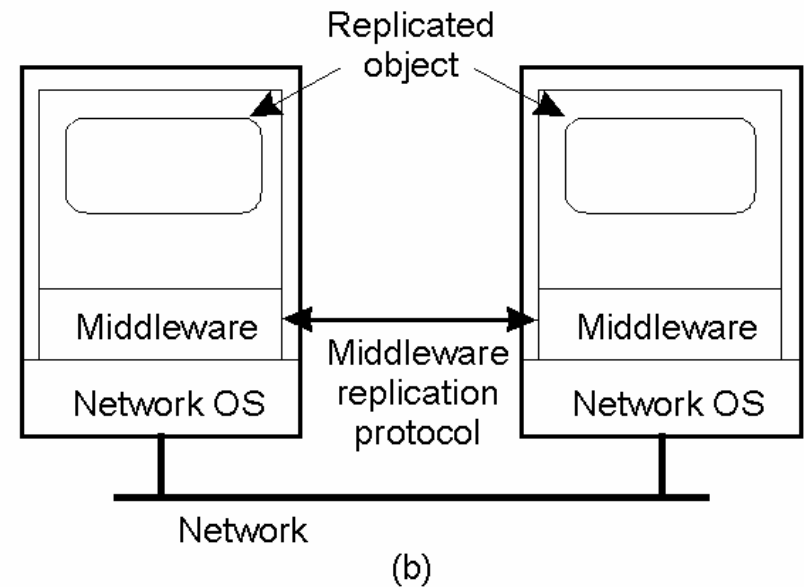
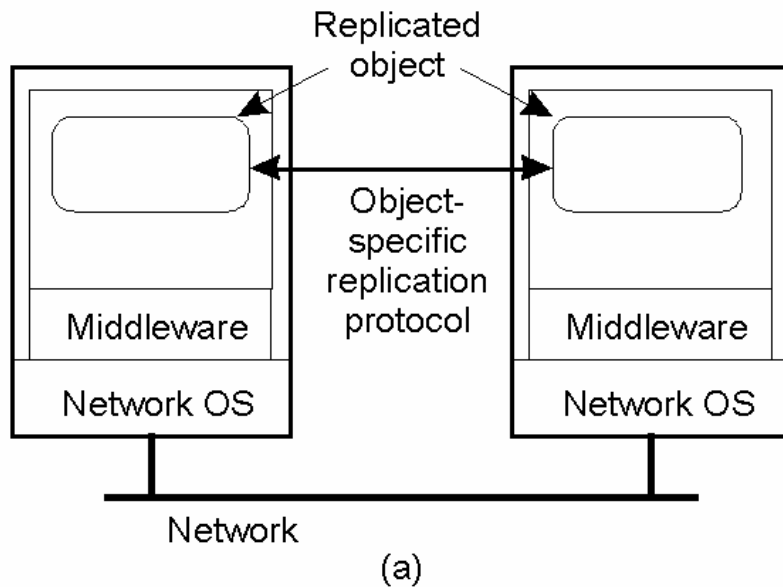
(a)



(b)


- a) A remote object capable of handling concurrent invocations on its own.
- b) A remote object for which an object adapter is required to handle concurrent invocations

# Object Replication (3)



- a) A distributed system for replication-aware distributed objects.
- b) A distributed system responsible for replica management

# Data-Centric Consistency Models



The general organization of a logical data store, physically distributed and replicated across multiple processes.

# Strict Consistency

P1:	W(x)a	
P2:		R(x)a

(a)

P1:	W(x)a	
P2:		R(x)NIL   R(x)a

(b)

Behavior of two processes, operating on the same data item.

- A strictly consistent store.
- A store that is not strictly consistent.

# Linearizability and Sequential Consistency (1)

P1:	W(x)a		
P2:	W(x)b		
P3:		R(x)b	R(x)a
P4:		R(x)b	R(x)a

(a)

P1:	W(x)a		
P2:	W(x)b		
P3:		R(x)b	R(x)a
P4:		R(x)a	R(x)b

(b)

(a) A sequentially consistent data store.

(b) A data store that is not sequentially consistent.

Process P1	Process P2	Process P3
x = 1; print (y, z);	y = 1; print (x, z);	z = 1; print (x, y);

Three concurrently executing processes.

# Linearizability and Sequential Consistency (2)

```
x = 1;  
print (y, z);  
y = 1;  
print (x, z);  
z = 1;  
print (x, y);
```

Prints: 001011

Signature:  
001011

(a)

```
x = 1;  
y = 1;  
print (x,z);  
print(y, z);  
z = 1;  
print (x, y);
```

Prints: 101011

Signature:  
101011

(b)

```
y = 1;  
z = 1;  
print (x, y);  
print (x, z);  
x = 1;  
print (y, z);
```

Prints: 010111

Signature:  
110101

(c)

```
y = 1;  
x = 1;  
z = 1;  
print (x, z);  
print (y, z);  
print (x, y);
```

Prints: 111111

Signature:  
111111

(d)

Four valid execution sequences for the processes of the previous slide. The vertical axis is time.



# Casual Consistency (1)

Necessary condition:

Writes that are potentially casually related must be seen by all processes in the same order. Concurrent writes may be seen in a different order on different machines.

P1:	W(x)a		W(x)c	
P2:		R(x)a	W(x)b	
P3:		R(x)a		R(x)c
P4:		R(x)a		R(x)c

This sequence is allowed with a casually-consistent store, but not with sequentially or strictly consistent store.

# Casual Consistency (2)

P1:	W(x)a		
P2:		R(x)a	W(x)b
P3:			R(x)b
P4:			R(x)a

(a)

P1:	W(x)a		
P2:		W(x)b	
P3:			R(x)b
P4:			R(x)a

(b)

- a) A violation of a casually-consistent store.
- b) A correct sequence of events in a casually-consistent store.

# FIFO Consistency (1)

Necessary Condition:

Writes done by a single process are seen by all other processes in the order in which they were issued, but writes from different processes may be seen in a different order by different processes.

P1: W(x)a

---

P2:            R(x)a    W(x)b    W(x)c

---

P3:                            R(x)b    R(x)a    R(x)c

---

P4:                            R(x)a    R(x)b    R(x)c

A valid sequence of events of FIFO consistency

# FIFO Consistency (2)

P1	P2	P3
x = 1;	x = 1;	y = 1;
<b>print (y, z);</b>	y = 1;	print (x, z);
y = 1;	<b>print(x, z);</b>	z = 1;
print(x, z);	print ( y, z);	<b>print (x, y);</b>
z = 1;	z = 1;	x = 1;
print (x, y);	print (x, y);	print (y, z);
Prints: 00	Prints: 10	Prints: 01
(a)	(b)	(c)

Statement execution as seen by the three processes from the previous slide. The statements in bold are the ones that generate the output shown.

# Weak Consistency (1)

Properties:

- Accesses to synchronization variables associated with a data store are sequentially consistent
- No operation on a synchronization variable is allowed to be performed until all previous writes have been completed everywhere
- No read or write operation on data items are allowed to be performed until all previous operations to synchronization variables have been performed.

# Weak Consistency (2)

P1:	$W(x)a$	$W(x)b$	$S$
P2:		$R(x)a$	$R(x)b$
P3:		$R(x)b$	$R(x)a$

(a)

P1:	$W(x)a$	$W(x)b$	$S$
P2:			$S \quad R(x)a$

(b)

- a) A valid sequence of events for weak consistency.
- b) An invalid sequence for weak consistency.

# Release Consistency

P1:	Acq(L)	W(x)a	W(x)b	Rel(L)	
P2:			Acq(L)	R(x)b	Rel(L)
P3:					R(x)a

A valid event sequence for release consistency.

## Rules:

- Before a read or write operation on shared data is performed, all previous acquires done by the process must have completed successfully.
- Before a release is allowed to be performed, all previous reads and writes by the process must have completed
- Accesses to synchronization variables are FIFO consistent (sequential consistency is not required).

# Entry Consistency (1)

## Conditions:

- An acquire access of a synchronization variable is not allowed to perform with respect to a process until all updates to the guarded shared data have been performed with respect to that process.
- Before an exclusive mode access to a synchronization variable by a process is allowed to perform with respect to that process, no other process may hold the synchronization variable, not even in nonexclusive mode.
- After an exclusive mode access to a synchronization variable has been performed, any other process's next nonexclusive mode access to that synchronization variable may not be performed until it has performed with respect to that variable's owner.



# Entry Consistency (2)

P1:	Acq(Lx)	W(x)a	Acq(Ly)	W(y)b	Rel(Lx)	Rel(Ly)
P2:					Acq(Lx)	R(x)a
P3:						R(y)b

A valid event sequence for entry consistency.

# Summary of Consistency Models

Consistency	Description
Strict	Absolute time ordering of all shared accesses matters.
Linearizability	All processes must see all shared accesses in the same order. Accesses are furthermore ordered according to a (nonunique) global timestamp
Sequential	All processes see all shared accesses in the same order. Accesses are not ordered in time
Causal	All processes see causally-related shared accesses in the same order.
FIFO	All processes see writes from each other in the order they were used. Writes from different processes may not always be seen in that order

(a)

Consistency	Description
Weak	Shared data can be counted on to be consistent only after a synchronization is done
Release	Shared data are made consistent when a critical region is exited
Entry	Shared data pertaining to a critical region are made consistent when a critical region is entered.

(b)

(a) Consistency models not using synchronization operations.

(b) Models with synchronization operations.

# Client-Centric Models

Eventual Consistency

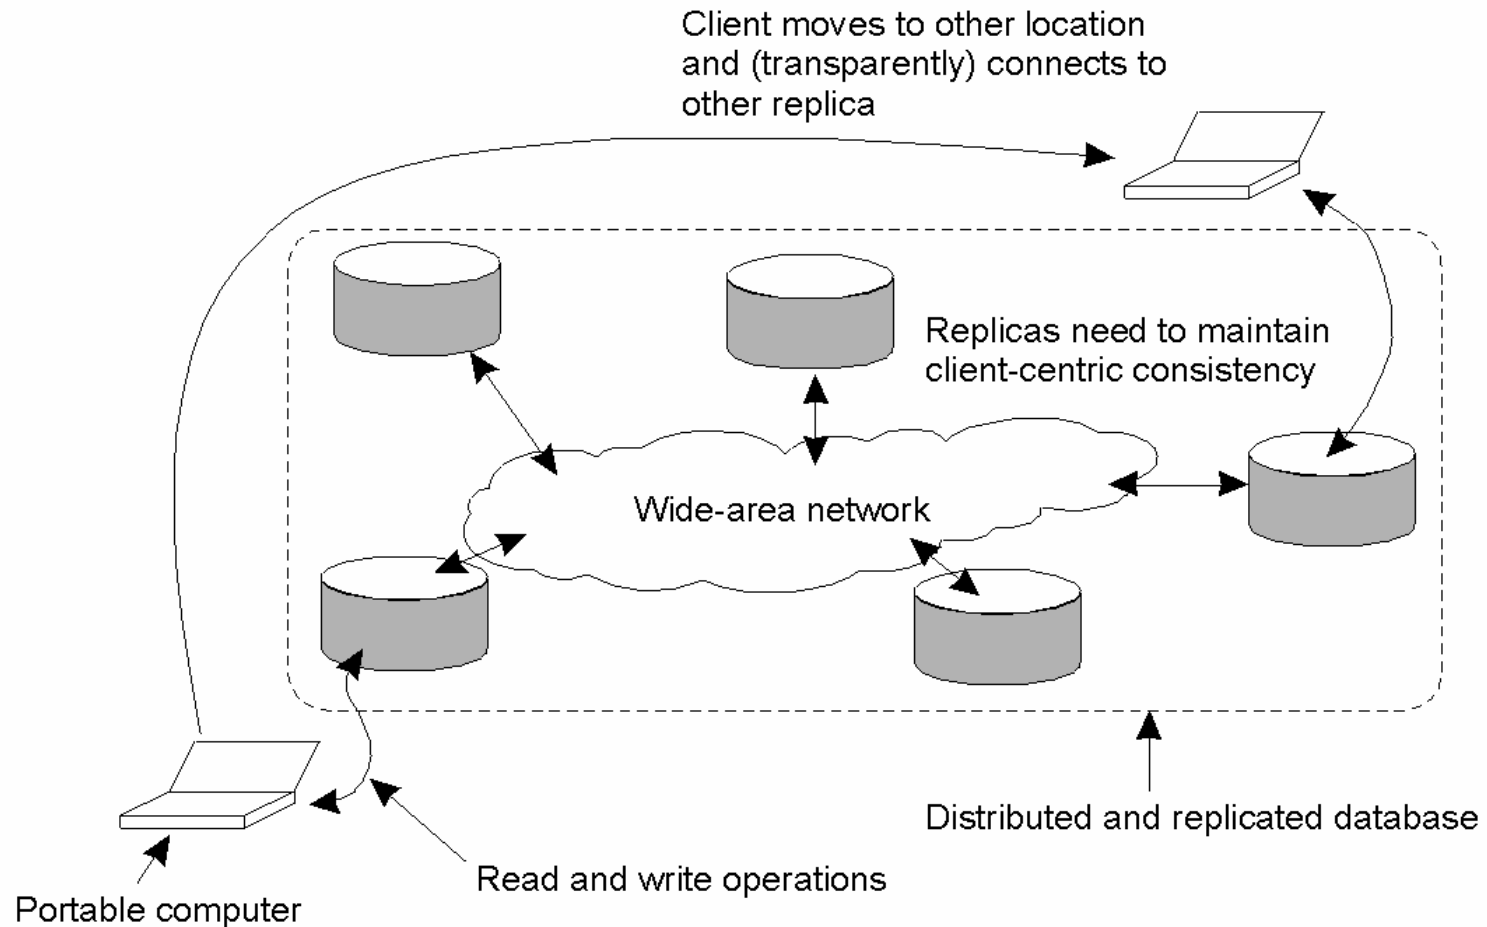
Monotonic Reads

Monotonic Writes

Read Your Writes

Writes Follow Reads

# Eventual Consistency



The principle of a mobile user accessing different replicas of a distributed database.

# Monotonic Reads

L1:	WS(x <sub>1</sub> )	R(x <sub>1</sub> )
<hr/>		
L2:	WS(x <sub>1</sub> ;x <sub>2</sub> )	R(x <sub>2</sub> )

(a)

L1:	WS(x <sub>1</sub> )	R(x <sub>1</sub> )
<hr/>		
L2:	WS(x <sub>2</sub> )	R(x <sub>2</sub> ) WS(x <sub>1</sub> ;x <sub>2</sub> )

(b)

The read operations performed by a single process  $P$  at two different local copies of the same data store.

- a) A monotonic-read consistent data store
- b) A data store that does not provide monotonic reads.

# Monotonic Writes

L1:	W(x <sub>1</sub> )	
<hr/>		
L2:	W(x <sub>1</sub> )	W(x <sub>2</sub> )

(a)



L1:	W(x <sub>1</sub> )	
<hr/>		
L2:		W(x <sub>2</sub> )

(b)

The write operations performed by a single process  $P$  at two different local copies of the same data store

- a) A monotonic-write consistent data store.
- b) A data store that does not provide monotonic-write consistency.

# Read Your Writes



- a) A data store that provides read-your-writes consistency.
- b) A data store that does not.

# Writes Follow Reads

L1:	WS(x <sub>1</sub> )	R(x <sub>1</sub> )
<hr/>		
L2:	WS(x <sub>1</sub> ;x <sub>2</sub> )	W(x <sub>2</sub> )

(a)

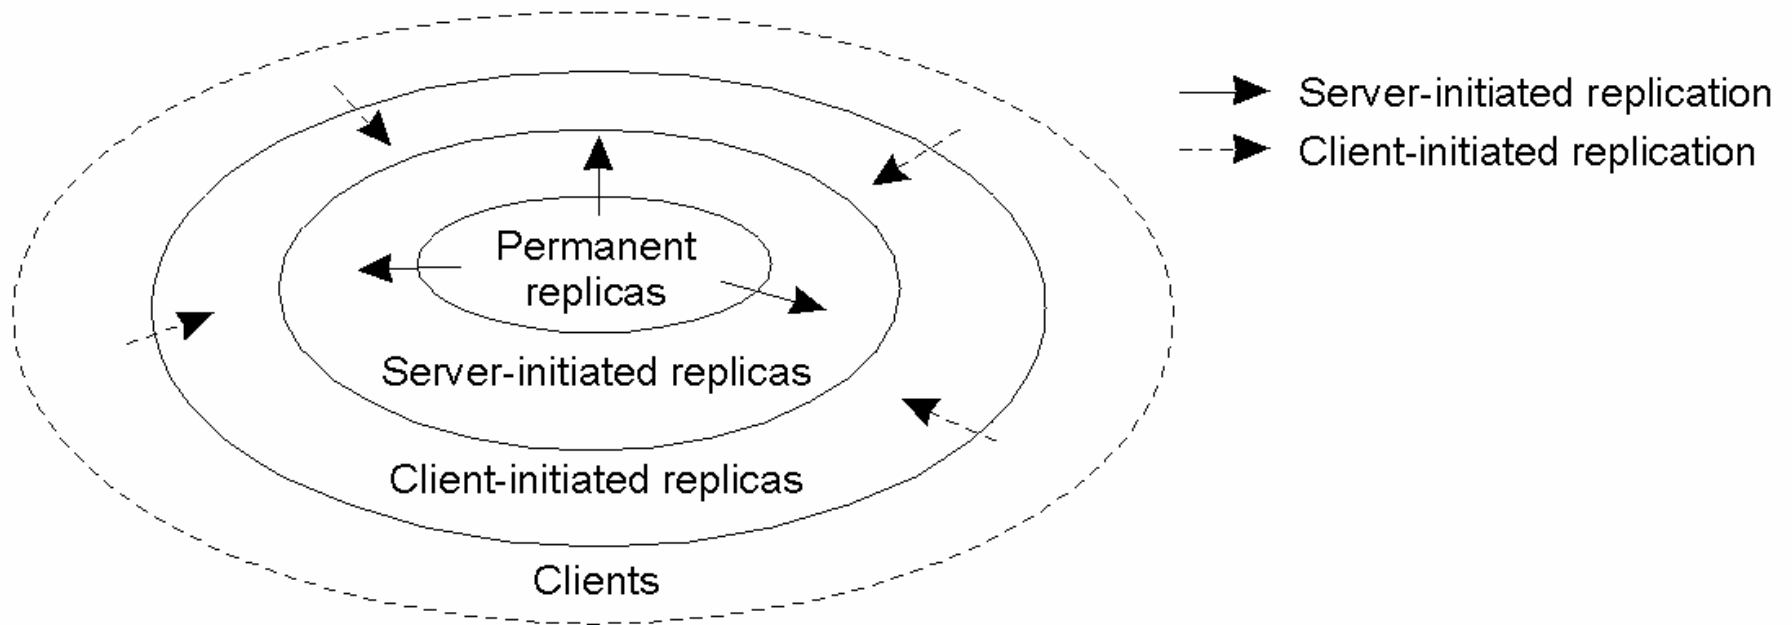
L1:	WS(x <sub>1</sub> )	R(x <sub>1</sub> )
<hr/>		
L2:	WS(x <sub>2</sub> )	W(x <sub>2</sub> )

(b)

- a) A writes-follow-reads consistent data store
- b) A data store that does not provide writes-follow-reads consistency




# Replica Placement



The logical organization of different kinds of copies of a data store into three concentric rings.

# Server-Initiated Replicas



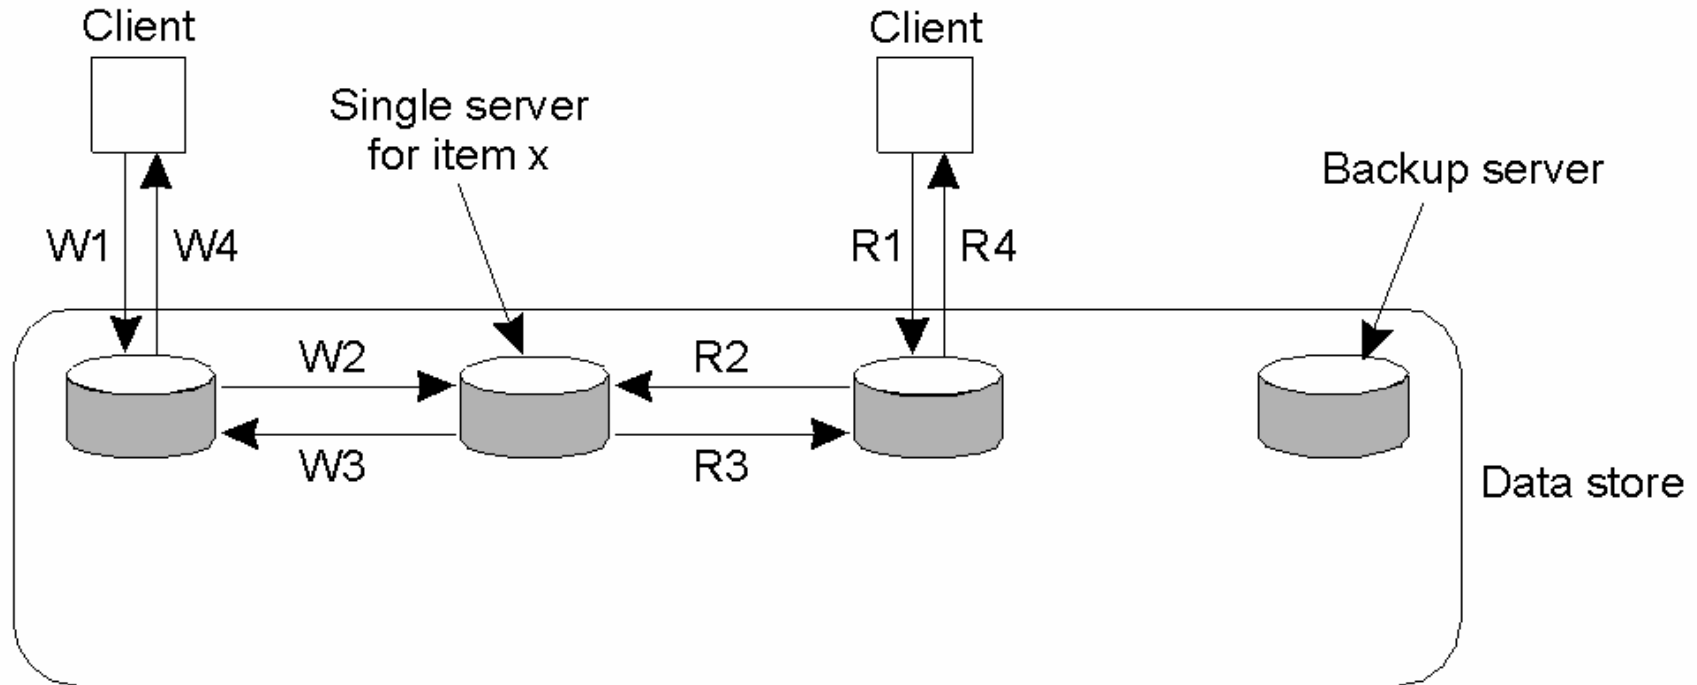
Counting access requests from different clients.

# Pull versus Push Protocols

Issue	Push-based	Pull-based
State of server	List of client replicas and caches	None
Messages sent	Update (and possibly fetch update later)	Poll and update
Response time at client	Immediate (or fetch-update time)	Fetch-update time

A comparison between push-based and pull-based protocols in the case of multiple client, single server systems.

# Remote-Write Protocols (1)

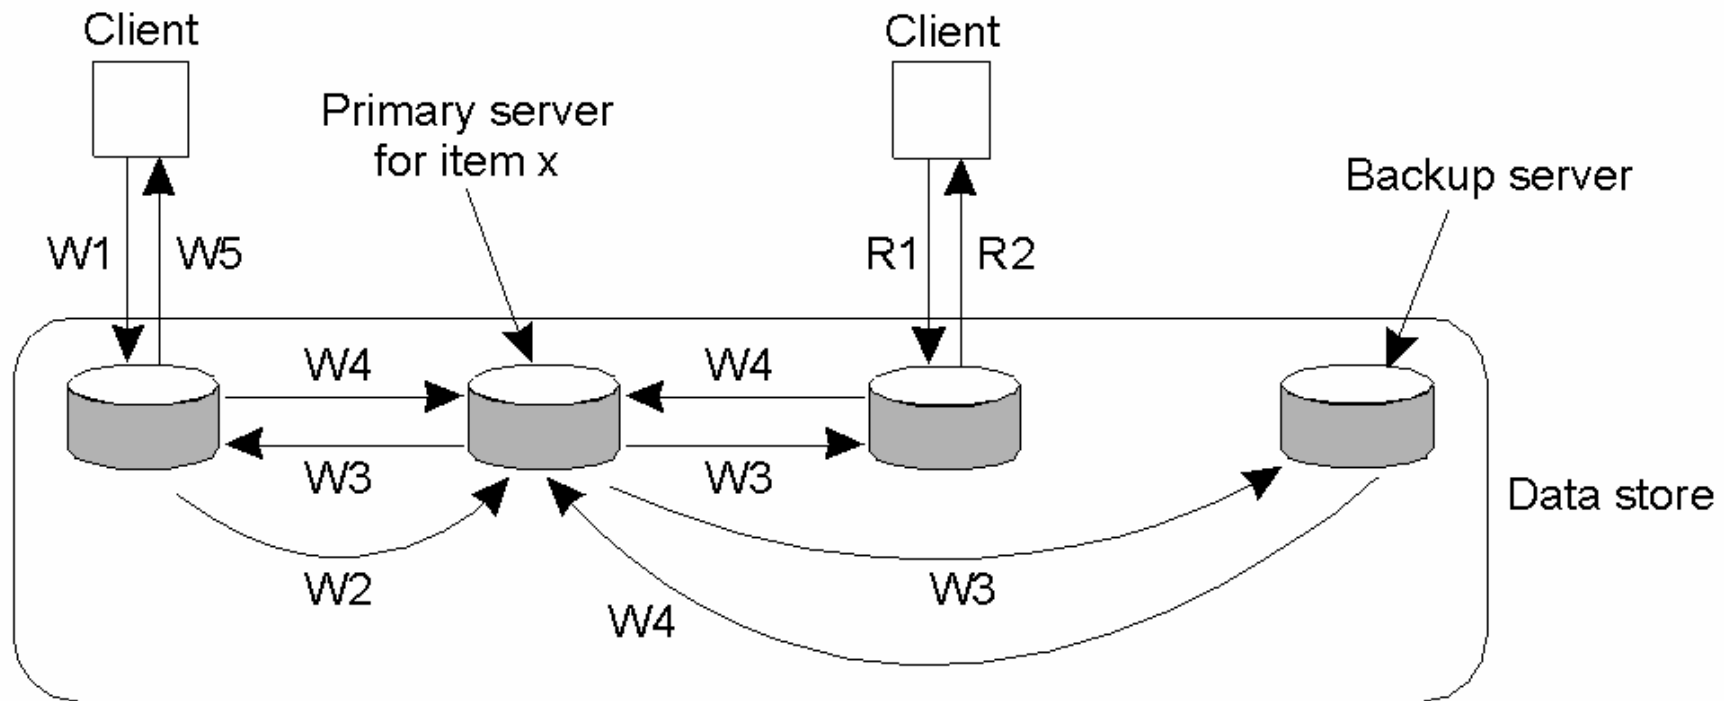


W1. Write request  
W2. Forward request to server for x  
W3. Acknowledge write completed  
W4. Acknowledge write completed

R1. Read request  
R2. Forward request to server for x  
R3. Return response  
R4. Return response

Primary-based remote-write protocol with a fixed server to which all read and write operations are forwarded.

# Remote-Write Protocols (2)

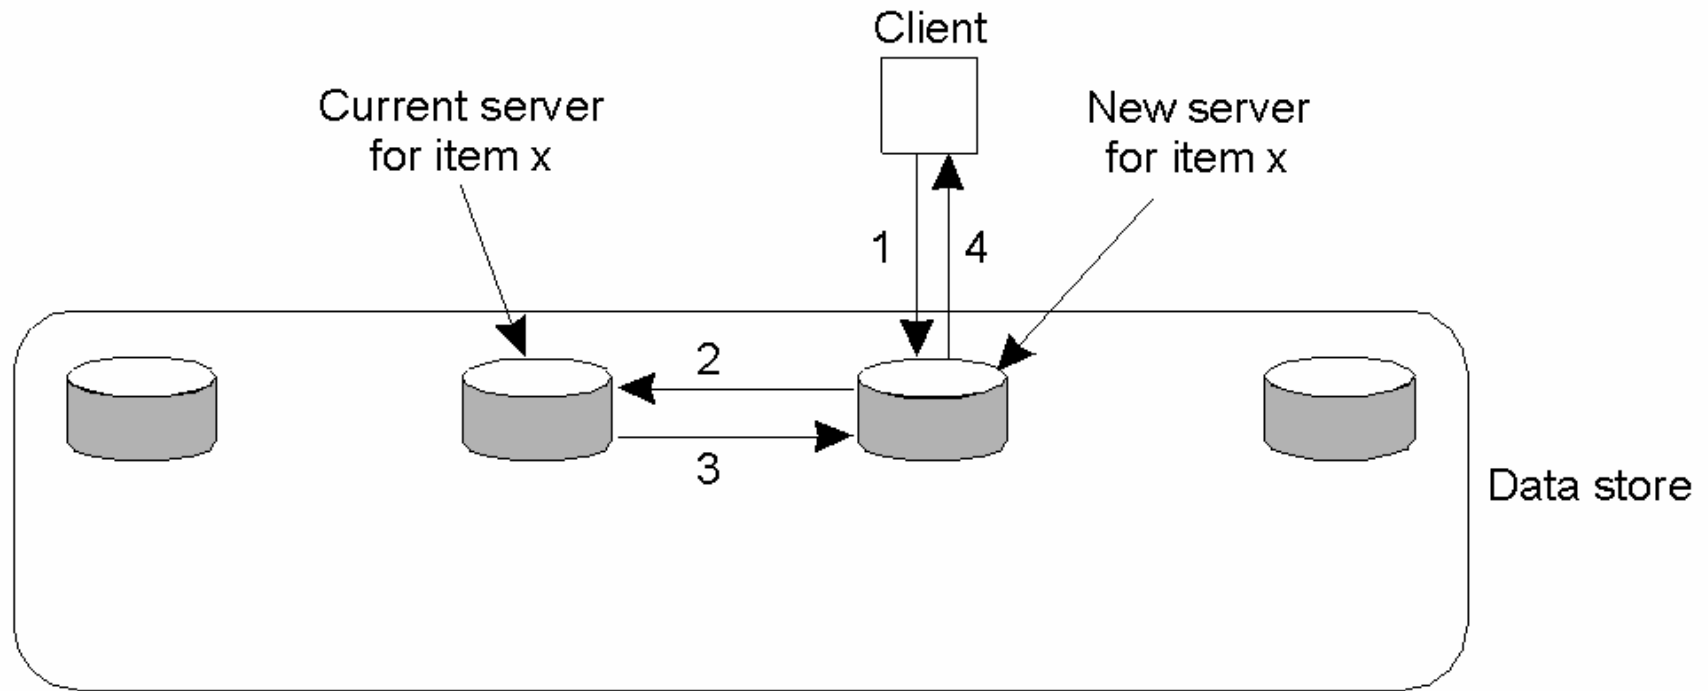


W1. Write request  
W2. Forward request to primary  
W3. Tell backups to update  
W4. Acknowledge update  
W5. Acknowledge write completed

R1. Read request  
R2. Response to read

The principle of primary-backup protocol.

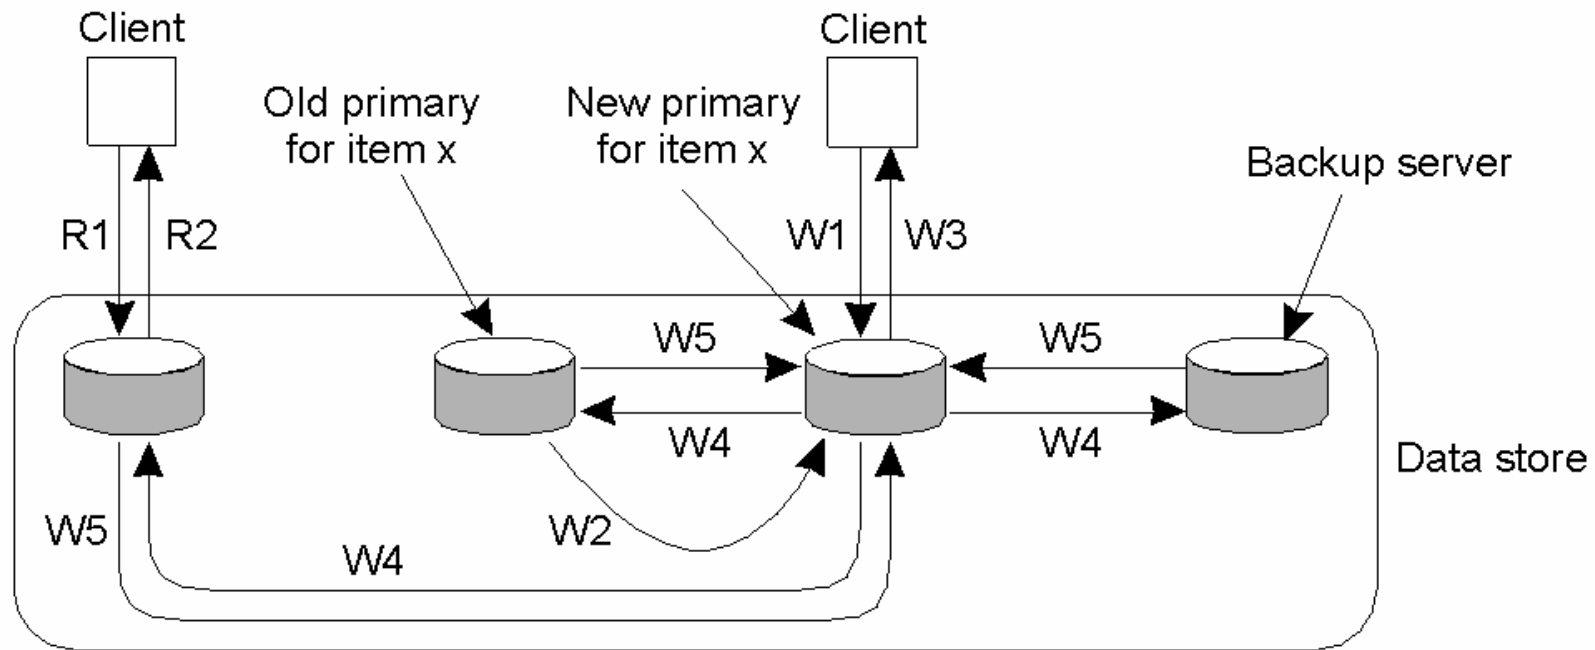
# Local-Write Protocols (1)



1. Read or write request
2. Forward request to current server for x
3. Move item x to client's server
4. Return result of operation on client's server

Primary-based local-write protocol in which a single copy is migrated between processes.

# Local-Write Protocols (2)

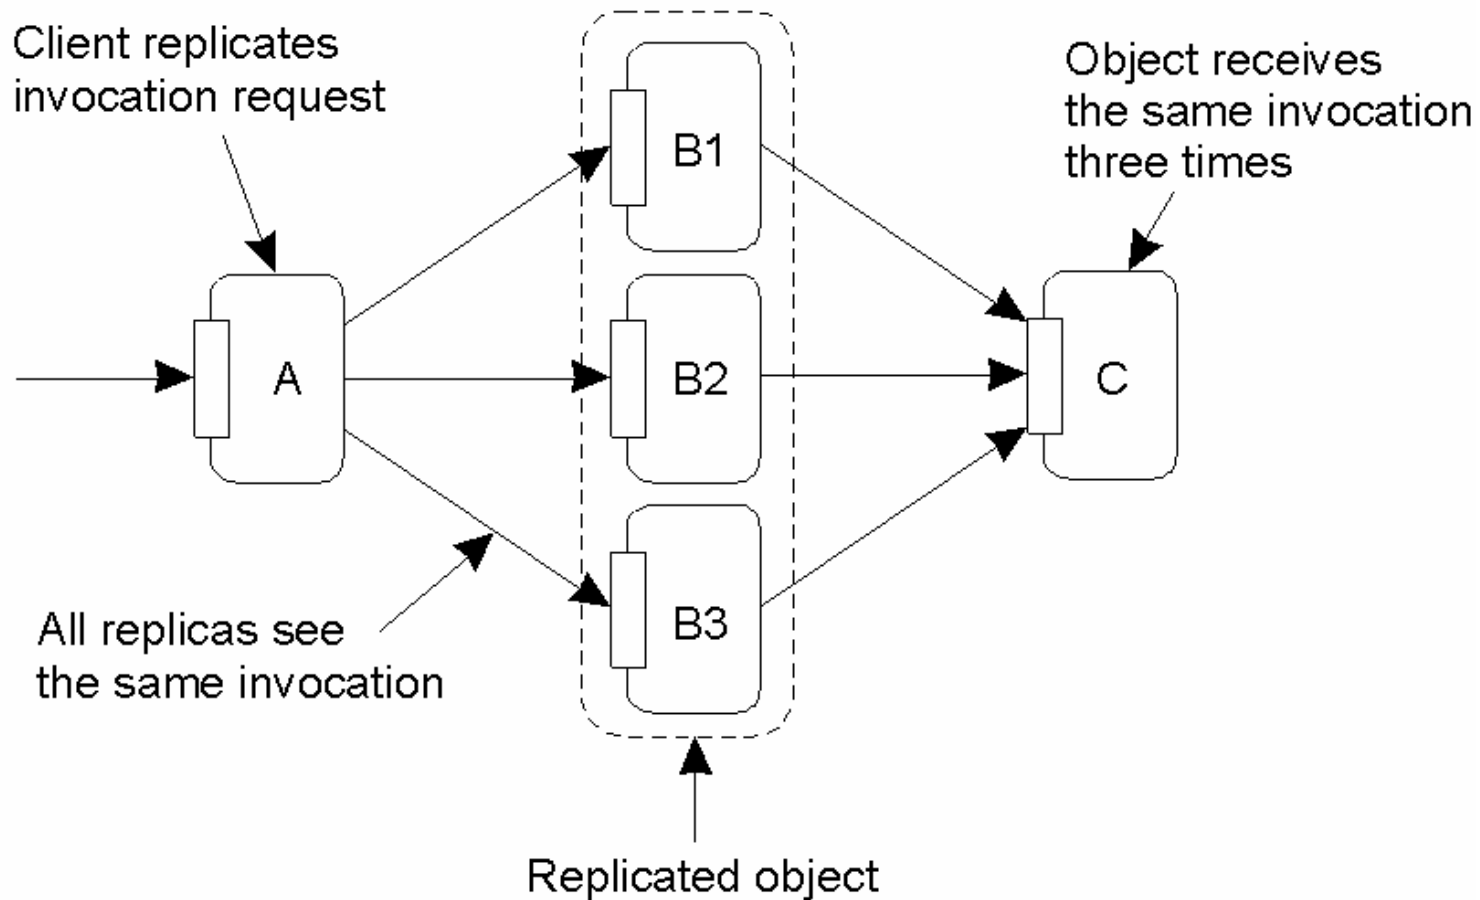


W1. Write request  
W2. Move item x to new primary  
W3. Acknowledge write completed  
W4. Tell backups to update  
W5. Acknowledge update

R1. Read request  
R2. Response to read

Primary-backup protocol in which the primary migrates to the process wanting to perform an update.

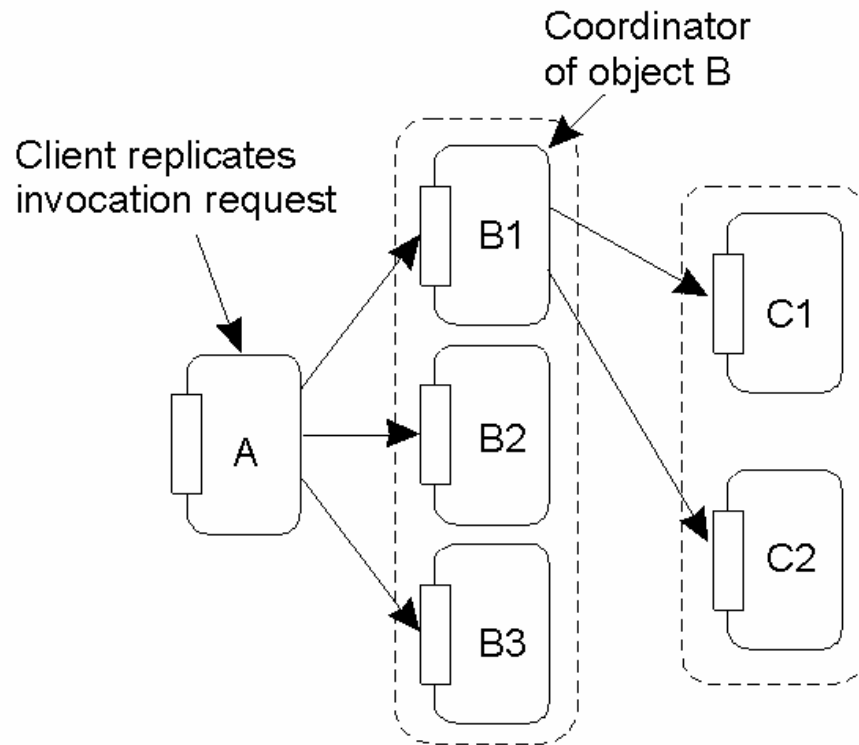
# Active Replication (1)




The problem of replicated invocations.



# Active Replication (2)



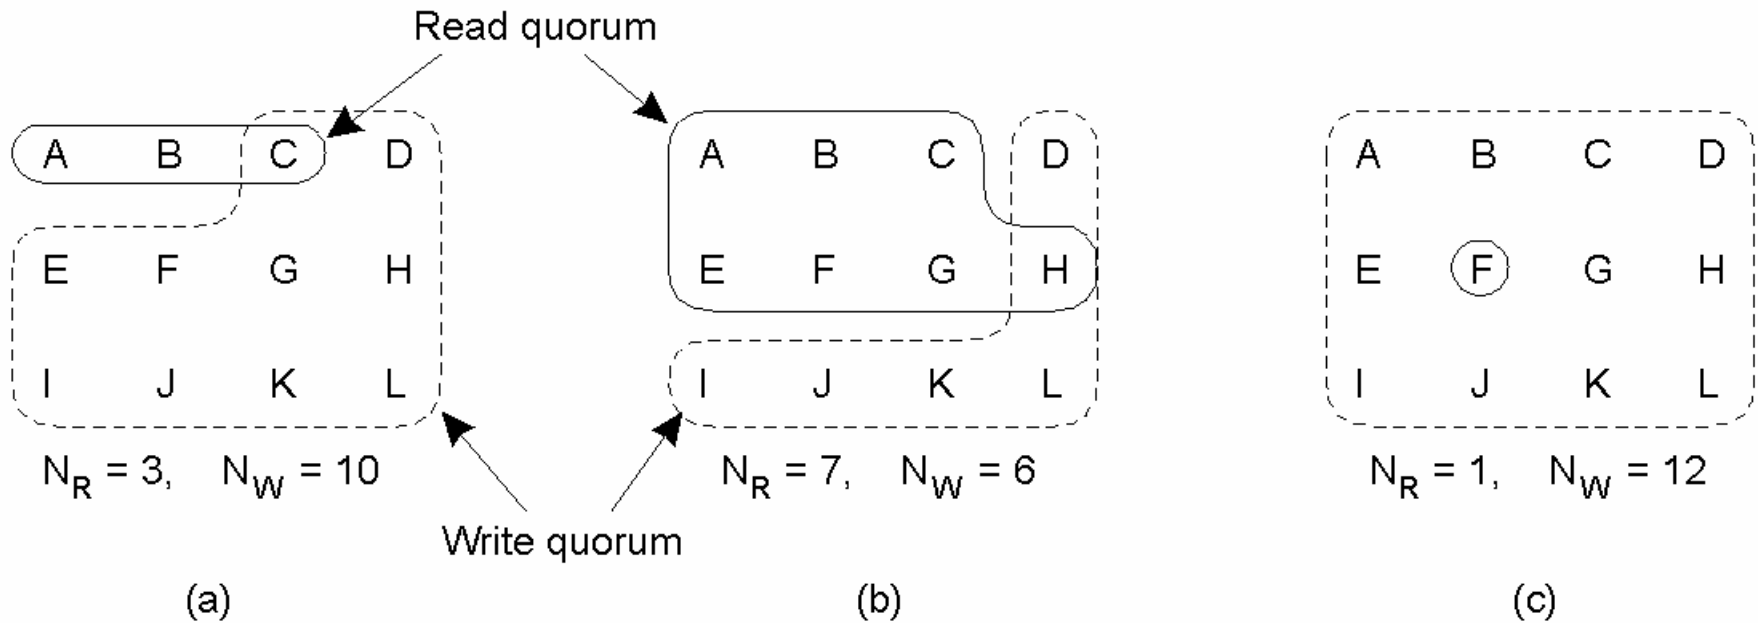
(a)



(b)

- a) Forwarding an invocation request from a replicated object.
- b) Returning a reply to a replicated object.

# Quorum-Based Protocols



Three examples of the voting algorithm:

- a) A correct choice of read and write set
- b) A choice that may lead to write-write conflicts
- c) A correct choice, known as ROWA (read one, write all)